

UNCLASSIFIED

Defense Technical Information Center  
Compilation Part Notice

ADP012349

TITLE: Knowledge Based Approach to OOTW Coalition Formation

DISTRIBUTION: Approved for public release, distribution unlimited

Availability: Hard copy only.

This paper is part of the following report:

TITLE: KSCO 2002: Second International Conference on Knowledge  
Systems for Coalition Operations

To order the complete compilation report, use: ADA402533

The component part is provided here to allow users access to individually authored sections of proceedings, annals, symposia, etc. However, the component should be considered within the context of the overall compilation report and not as a stand-alone technical report.

The following component part numbers comprise the compilation report:

ADP012330 thru ADP012354

UNCLASSIFIED

# Knowledge Based Approach to OOTW Coalition Formation

Michal Pěchouček, Vladimír Mařík, Jaroslav Bárta

Gerstner Laboratory, Department of Cybernetics,  
Czech Technical University,  
Technická 2, 166 27 Prague, Czech Republic  
email: {pechouc|marik|barta}@labe.felk.cvut.cz

**Abstract.** The task of planning humanitarian relief operations within a high number of hardly collaborating and vaguely linked non-governmental organizations is a challenging problem. We suggest an alternative knowledge based approach to the coalition formation problem for humanitarian and peace-keeping missions. Owing to the very special nature of this domain, where the agents representing individual organisations may eventually agree to collaborate, but are very often reluctant to share their knowledge and resources, we tried to reduce the problem complexity by splitting the community of agents into alliances. We combined classical negotiation mechanisms with the acquaintance models and social knowledge techniques in order to reduce the communication traffic and to keep the privacy of knowledge. Experimental results are discussed in the paper.

## 1 Introduction

The application domain of this coalition formation research belongs to the area of **war avoidance operations** such as peace-keeping, peace-enforcing, non-combatant evacuation or disaster relief operations. Unlike in classical war operations, where the technology of decision making is strictly hierarchical, **operations other than war** (OOTW) are very likely to be based on cooperation of a number of different, quasi-volunteered, vaguely organized groups of people, non-governmental organizations (NGO's), institutions providing humanitarian aid, but also army troops and official governmental initiatives.

**Collaborative**, unlike hierarchical, approach to operation planning allows greater deal of flexibility and dynamics in grouping optimal parties playing an active role in the operation. New entities shall be free to join autonomously and involve themselves in planning with respect to their capabilities. Therefore any organization framework must be essentially "open". OOTW have, according to (Walker, 1999), multiple perspective on plan evaluation as there does not need to be one shared goal or a single metrics of the operation (such as political, economical, humanitarian). From the same reason, the goals of entities involved in a possible coalition may be in conflict. Even if the community members share the same goal, it can be easily misunderstood due to different cultural backgrounds.

The main reason why we can hardly plan operations involving different NGO's by a central authority results from their **reluctance to provide information** about their intentions, goals and resources. Consequently, besides difficulties related to planning and negotiation we have to face the problems how to assure sharing the detailed information. Many institutions will be ready to share resources and information within some well specified community, whereas they will refuse to register their full capabilities and plans with a central planning system and to follow centralized commands. They may agree to participate in executing a plan, in forming of which they played an active role. In our interpretation, an agent is a complex, organized entity (representing a NGO, humanitarian organization, army troop, etc.) playing an active role in the OOTW planning. A multi-agent system consists of a number of agents that group themselves in various, temporary coalitions (each solving a specific mission/part of the mission).

The main ambition of our research has been to analyze the problem of OOTW coalition formation and to propose a novel approach that would (i) make the coalition formation process simpler in comparison to the "classical" methods, and thus more efficient and (ii) at the same time maintain confidentiality of the private information. In our case, we decided to sacrifice the total optimality of the formed coalitions as we found this is not the most important aspect in the OOTW planning. We have suggested a concept of alliances – a set of agents that agreed to share some of their private information and to cooperate eventually. The coalition formation complexity is reduced by splitting the whole community of agents into disjunctive subsets (alliances) and by the attempts to create a coalition preferably within the single alliance. Social knowledge stored in the acquaintance models of individual agents has been widely explored in order to:

- **minimize required communication traffic** which influences the problem solving efficiency,
- **keep the quality of the coalition** that resulted from the coalition formation process operation 'reasonably good' – the quality has been measured by the humanitarian relief aid deliver time and by how much the coalition covers the request (in percent),
- **minimize the loss of agents' semiprivate information** when negotiating the mission – i.e. revealing the information about services the agent may provide, its status and intention in the minimum extent, and
- **minimize the amount of shared information** – information that possible coalition leaders know about other agents and use it in order to plan an optimal mission.

The developed approach has been tested on the CPlanT multi-agent system implementation.

## 2 CPlanT System Architecture

CPlanT is a multi-agent system for planning humanitarian relief operations where any agent can initiate the planning process. Classical negotiation algorithms such as contract net protocol (CNP) are used in combination with the acquaintance models techniques (Smith, 1980; Mařík, 2001). The CPlanT architecture consists of several specific classes of agents:

**Resource Agents (R-agents)** represent the in-place resources that are inevitable for delivering humanitarian aid, such as roads, airports. Unlike the below-defined H-agents, the R-agents are regarded as passive and they do not initiate any kind of humanitarian effort.

**In-need Agents (In-agents)** represent the centers of conflict that call for help (e.g. cities, villages, etc.).

**Humanitarian Agents (H-agents)** represent the participating humanitarian agencies. Like the R-agents, the H-agents contribute to humanitarian aid missions. Therefore, one may regard the H-agent as a subclass of R-agents. However the H-agents are proactive and they can initiate coalition formation process.

In this paper, we will investigate coalition formation among the H-agents.

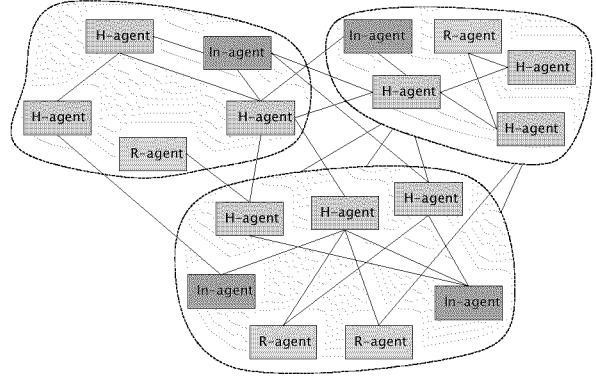


Figure 1 – CPlanT Multi-Agent Architecture

## 3 Knowledge Architecture

### 3.1 Agent's Neighborhood

Each H-agent may participate in one alliance of 'friendly' agents and at the same time it may be actively involved in several coalitions of agents cooperating in fulfilling specific shared tasks. Computational and communication complexity of forming such a coalition depends on the amount of pre-prepared information the agents administer one about the other and on sophistication of the agents' capability to reason about the other agents' resources, plans and intentions. The agents can allow others to reason about them and at the same time they can reason differently about the agents that belong to their different scopes of reasoning – neighborhood. Therefore, we distinguish among several types of agents' neighborhoods:

- $\alpha(A)$  – agent's **total neighborhood**, a set of all agents that the agent  $A$  is aware of, (e.g. knows about their existence and is able to communicate with them)
- $\mu(A)$  – agent's **social (monitoring) neighborhood** that is a set of agents, which the agent  $A$  keeps specific information about (e.g. services they provide, status, load, etc.). This neighborhood consists of the set of the agents that the agent  $A$  reasons about –  $\mu^+(A)$  and the set the agents that reason about the agent  $A$  –  $\mu^-(A)$ . Therefore

$$\forall B \in \mu^-(A): A \in \mu^+(B).$$

- $\varepsilon(A)$  – agent's **cooperation neighborhood** that is a set of agents jointly collaborating (or committed to collaboration) in achieving one or more shared goals.

### 3.2 Knowledge Sharing

In order to reason one about the other, the agents must share some of their knowledge. Let us introduce the operator  $(\text{Bel } A \ \varphi)$  that expresses the agent's  $A$  awareness of the formula  $\varphi$  being true (Wooldirdge 2000). We say that the agent  $A_0$  intentionally shares its knowledge  $\mathbf{K}(A_0)$  with a set of agents  $\delta(A_0) \subseteq \Theta$  provided that:

$$\mathbf{K}(A_0) = \{\varphi\} : \forall \varphi \in \mathbf{K}(A_0) : \forall A_i \in \delta(A_0) : (\text{Bel } A_i \ \varphi) \wedge \forall B_i \notin \{\delta(A_0) \cup \{A_0\}\} : (\text{Bel } A_0 \neg (\text{Bel } B_i \ \varphi)).$$

From the previous follows, that if an agent  $B$  knows some of the shared information without the agent  $A_0$  being aware of this fact, the agent  $B$  is not regarded as a member of the  $\delta(A_0)$  set of agents, representing  $A_0$ 's knowledge sharing neighborhood. According to this classification, we suggest three levels of the H-agent's knowledge sharing:

**Public Knowledge** is shared within the entire multi-agent community. If it is assumed that all the agents know one about the other (i.e.  $\forall A, A \in \Theta : \alpha(A) = \Theta$ ), public knowledge  $\mathbf{K}_p(A_0)$  of an agent  $A_0$  is defined as

$$\mathbf{K}_p(A_0) = \mathbf{K}(A_0) \text{ where } \delta(A_0) = \alpha(A_0).$$

This class of knowledge is freely accessible within the community. As public knowledge we understand the agent's name, the type of the organization the agent represents, the general objectives of the agent's activity, the country

where the agent is registered, agent's human-human contact (telephone, fax number, email), the human-agent type of contact (http address), the agent-agent type of contact (the IP address, incoming port, ACL) and, finally, available services.

**Semi-Private Knowledge** is shared within agents' social neighborhoods. Semi-private knowledge  $K_s(A_0)$  of an agent  $A_0$  is defined as

$$K_s(A_0) = K(A_0) \text{ where } \delta(A_0) = \mu(A_0).$$

As in the OOTW domain, we do not assume the knowledge to be shared within the overlapping alliances, we will require the social neighborhood to have the following property:  $\forall A \in \Theta : \mu^-(A) = \mu^+(A) = \mu(A)$ . Members of a social neighborhood share information about availability of their resources.

**Private Knowledge** is owned and administered by the agent itself. Private knowledge  $K_p(A_0)$  of an agent  $A_0$  is defined as

$$K_p(A_0) = K(A_0) \text{ where } \delta(A_0) = \{\},$$

An important type of private knowledge includes agent's collaboration preferences, alliance restrictions, coalition leader restrictions and possible next restrictions, but also agent's planning and scheduling algorithms.

### 3.3 Alliance, Coalition, Team Action Plan

In the subject domain, we will understand as the multi-agent community  $\Theta$  the whole collection of agents participating in the above-described OOTW (quasi-volunteered, vaguely organized groups of people, non-governmental organizations, institutions providing humanitarian aid, army troops or official governmental initiatives). We will introduce the concept of an **alliance** as a collection of agents that share information about their resources and all agree to form possible coalitions. The alliance is regarded as a long-term cooperation agreement among the agents. Members of an alliance will all belong to one others' social neighborhood. Provided that we assume that each agent belongs also to its own social neighborhood –  $\forall A \in \Theta : A \in \mu(A)$ , we define the alliance as follows:

An **alliance** is a set of agents  $\kappa$ , so that  $\forall A \in \Theta : \exists \kappa : A \in \kappa \wedge \forall A_i \in \kappa : \kappa = \mu(A_i)$ .

A singleton agent is regarded as an alliance with just one member. From the requirements for the reciprocal knowledge sharing within an alliance follows that

$$\forall A \in \kappa : \kappa = \mu(A).$$

Therefore, an important property of an alliance is that it cannot overlap with another alliance:

$$\forall \kappa_1, \kappa_2 \subseteq \Theta : (\exists A : A \in \kappa_1 \wedge A \in \kappa_2) \Rightarrow \kappa_1 \equiv \kappa_2.$$

Let us define a **coalition** as a set of agents, which agreed to fulfill a single, well-specified goal. Coalition members committed themselves to collaborate on the within-coalition-shared goal. Under the assumption  $\forall A \in \Theta : A \in \varepsilon(A)$  we define coalition as follows:

A **coalition** is a set of agents  $\chi$ , so that  $\forall \chi(\tau) \subseteq \Theta : \forall A \in \chi(\tau) : \chi(\tau) \subseteq \varepsilon(A)$ .

Let us introduce a set  $\varepsilon(A, \tau)$  that is an agent collaboration neighborhood with respect to a single shared goal  $\tau$ . Then

$$\varepsilon(A) = \bigcup_{\tau} \varepsilon(A, \tau), \text{ and } \forall \chi(\tau) \subseteq \Theta : \forall A \in \chi(\tau) : \chi(\tau) \subseteq \varepsilon(A, \tau).$$

A coalition, unlike an alliance, is usually regarded as a short-term agreement between collaborative agents. As we will see in Section 6, it is better for a coalition to be a subset of one alliance, but it is not an inevitable condition. A coalition can consist of agents who are members of different alliances.

Another term that we have to introduce is a **team action plan**. In planning humanitarian relief operations, similarly as in the case of any other collaborative action planning, the agents must agree on how they will achieve the goal  $\tau$ . The team action plan is thus a decomposition of a goal  $\tau$  into a set of tasks  $\{\tau_i\}$ . The tasks will be delegated within the coalition members. Apart from the responsible agent, each task shall be denoted by its due time, start time and price. Provided that an agent  $A_i$  is responsible for implementing the task  $\tau_i$  (where  $\tau = \{\tau_i\}$ ) in time  $\text{due}(\tau_i)$ , starting at  $\text{start}(\tau_i)$  for the price  $\text{price}(\tau_i)$ , we define the team action plan as follows:

A **team action plan**  $\pi(\tau)$  is as a set  $\pi(\tau) = \{\langle \tau_i, A_i, \text{start}(\tau_i), \text{due}(\tau_i), \text{price}(\tau_i) \rangle\}$ .

We say that the team action plan  $\pi(\tau)$  is **correct** if all the collaborators  $A_j$  are able to implement the task  $\tau_i$  in the given time and for the given price. The team action plan  $\pi(\tau)$  is **accepted** if all agents  $A_j$  get committed to implementing the task  $\tau_i$  in the given time and for the given price. Similarly, we say about the goal  $\tau$  to be **achievable**, if there exists such  $\pi(\tau)$  that is correct. The goal  $\tau$  is said to be **planned**, if there exists  $\pi(\tau)$  that is accepted. Obviously, there is an important relation between the team action plan and the coalition. We say that a coalition  $\chi(\tau)$  achieves a goal  $\tau$  by implementing a team action plan  $\pi(\tau)$  if and only if  $\chi(\tau) = \{A_j\}$  and  $\pi(\tau)$  is correct.

### 3.4 Disclosure of Private and Semi-Private Knowledge

Measuring the loss of information, that the agents may want to keep private, is an uneasy task. The revealed piece of information has got different value to agents with different meta-reasoning capabilities (Pěchouček & Norrie, 2000). In order to vaguely categorize various types of information leaks, let us distinguish between strong and weak leaks.

- **Strong information disclosure:** If an agent loses some type of private (or semi-private) knowledge in the strong sense, it does so as a side effect of some proactive step (such as sending a request).
- **Weak information disclosure:** If an agent loses the private knowledge in the weak sense, it deliberately discloses some piece of its knowledge to other agents (e.g. when sending an `inform`-type message).

Each agent undertakes the weak loss of some of its knowledge when forming an alliance. At this moment the agent's semi-private knowledge gets disclosed within the alliance members. In our system, the agent loses some of its private knowledge in the strong sense, if it communicates with an agent which is outside of its alliance. Once the agent  $A_1$  from an alliance  $\kappa_1$  sends a request for a service  $\tau$  to the agent  $A_2$  from the alliance  $\kappa_2$ , the agent  $A_1$  reveals the information about the **intent** (e.g.  $A_1$  does something that requires  $\tau$ ) and information about agent's  $A_1$  capabilities (e.g.  $A_1$  cannot do  $\tau$ ). At the same time, a proposal for collaboration from the agent  $A_2$  reveals the information about agent's  $A_2$  capabilities (such as  $A_2$  can implement  $\tau$  in time  $t_1$ ). However, this type of knowledge disclosure has been reduced as the agent  $A_2$  acts on behalf of the entire alliance. Therefore, if  $A_2$  offers some services that are not used at the end, there is a loss of information about capabilities of the entire alliance (and not of the agent  $A_2$  itself).

## 4 Agents' Acquaintance Model

Let us very briefly introduce the concept of agent's social intelligence and acquaintance models. Apart from its **problem-solving knowledge** that guides agent's autonomous local decision making processes (such as coalition formation, or team action planning), the agents usually exploit **social knowledge** that expresses the other agent's behavioral patterns, their capabilities, load, experiences, resources, commitments, knowledge describing conversations or negotiation scenarios (Mařík et. al., 2001). This knowledge is usually stored separately from the agents' computational core – in an agent's **acquaintance model**. There have been investigated several acquaintance models previously. Based on the *tri-base acquaintance model* (Pěchouček et. al., 2001), the social knowledge in CPlanT is organized in four separate knowledge structures:

- **community-base** (Com-BB) – which is a collection of the community members' public knowledge  

$$\text{Com-BB}(A_0) = \{K_p(A_i)\} \text{ for } \forall A_i \in \alpha(A_0)$$
- **self-belief-base** (Self-BB) – where the agent's reflective knowledge about itself is located; here the agent stores its public knowledge that is accessible to anyone, its semi-private knowledge that is shared within the alliance and its private knowledge that is not shared by anyone,  

$$\text{Self-BB}(A_0) = \{K_p(A_0)\}, \{K_s(A_0)\}, \{K_r(A_0)\}$$
- **social-belief-base** (Soc-BB) – where the agent stores the semi-private knowledge of its peer alliance members,  

$$\text{Soc-BB}(A_0) = \{K_s(A_i)\} \text{ for } \forall A_i \in \mu(A_0)$$
- **coalition-base** (Coal-BB) – which is a dynamic collection of the peer coalition members, the past and possible future coalitions as much as permanent coalition-formation rules<sup>1</sup>.




Figure 2 – Structure of the CPlanT Acquaintance Model

Exploitation of the acquaintance model reduces communication traffic required for collaborative activity planning. In principle, the social knowledge substantially reduces the set of agents (ideally to one) that will be requested by the coordinating agent in the CNP process (Smith, 1980). An important flaw of this approach is rooted in high requirements for the social model maintenance. The social knowledge maintenance may be driven either by the owner of the acquaintance model (the coordinator) or by those which are represented in the model – hence service providers (collaborators). We refer to the former strategy as the **requestor-driven** knowledge maintenance and to the latter strategy as the **provider-driven** knowledge maintenance. As an example of a **requestor-driven** strategy

<sup>1</sup> The coalition-formation rules are instances of the agent's problem-solving knowledge, while the information about the coalition members, past and future coalitions are instances of the social knowledge. Therefore the coalition base belongs in part to the acquaintance model and to the agent's body

there is the concept of **periodical revisions** (Mařík et.al., 2000) where the knowledge owner periodically checks consistency of the model with the potential collaborators. In other systems, there has been a **cooperation trader** (Cao et. al., 1997) type of agent, which was in charge of maintaining the agents social knowledge. As explained in Section 164 we have adopted the **provider-driven** knowledge maintenance in CPlanT.

Self-Belief Base		
public knowledge:	Semi-private knowledge:	Private knowledge
<b>Port:</b> 1500 <b>ip_address:</b> "147.32.86.167" <b>Country:</b> suffer terra <b>City:</b> north port <b>Type:</b> Religious <b>Ontologies:</b> fipa-am, cplant-ontology	Food: 3000 Nurses: 50 Trucks: 13	Alliance restrictions: ("country", "Suffer Terra") Leader restrictions: ("type", "Military"). City restrictions: ("muslim", 50) Cooperates with: ("type", "government")
Social belief base		
	<b>Agent:</b> ST Police <b>Agent:</b> Christian STHO	Armed-people:30 Food: 3500 Nurses: 22 Clothing: 280 Medical-people: 12
Community belief base		
	<b>Agent:</b> Suffer Terra Government <b>Agent:</b> Christian STHO	Suffer Terra Government@iioip://147.32.84.131:2188/Suffer Terra Government Type: Government Services: Food, Civil-material, Medical-material, Clothing Ontologies: FIPA-AGENT-MANAGEMENT, MAP-ONTOLOGY, PORT-ONTOLOGY, CPLANT, ALLIANCE Languages: SL1, KIF, State: ACTIVE Country: Suffer Terra, City: Suffer Town Christian Suffer Terra Humanitarian Organization@iioip://147.32.84.131:2210/Chr ST Humanitarian Organization Type: Religious Services: Food, Clothing, Medical-people, Nurses, Medical-material Ontologies: FIPA-AGENT-MANAGEMENT, MAP-ONTOLOGY, PORT-ONTOLOGY, CPLANT, ALLIANCE Languages: SL1, KIF, State: ACTIVE Country: Suffer Terra, City: North Port
Coalition Base		
Rules	(VOLCANIC-AVERAGE-SMALL-TOWN → Time: 220 (Requirements: Medical-material 60, Food 1500, Civil-material 30000, Medical-people 16, Civil-people 27, Nurses 19) ...)	
Coalitions	(coalition ( <b>Members:</b> Suffer Terra Government, Suffer Terra Police, Christian Suffer Terra Humanitarian Organization) ( <b>Services:</b> Food, Civil-material, Medical-material, Clothing, Military-people, Food, Clothing, Medical-people, Nurses) ( <b>Price-for-coordination:</b> 5)) (planned-coalition ( Task name: Suffer-Town-24-1-2002/17-49-53.1 (Coalition members: Suffer Terra Government, Suffer Terra Police, Christian Suffer Terra Humanitarian Organization) (Coalition leader: Christian Suffer Terra Humanitarian Organization (Disaster: Volcanic, Degree: 1, (Allocations: Civil-material, 80000, Allocation Time: 350 Food, 80000, Allocation Time: 350 Medical-material, 80000, Allocation Time: 350)) ...	

Table 1 – Instance of an H-agent's acquaintance model

## 5 Inter-Agent Communication

Before explaining the lifecycle of the system let us comment the main communication techniques that have been used in CPlanT: central communication agent, contract net protocol, and acquaintance models. We have tried to minimize the role of the central communication component, as it is an important communication bottleneck of the system operation and a center where the agents' private knowledge may be sniffed (see Section 6).

### 5.1 Contract Net Protocol

The CPlanT implementation relied heavily on the **contract net protocol** (CNP) negotiation scenario (Smith, 1980). Any agent can initiate the coalition forming process (hereafter we refer to this agent as a coalition **coordinator**) by requesting some agents in the community (**collaborators**) for specific services. Upon receiving proposals for collaboration, the coordinator carries out a computational process by which it selects the best possible collaborator(s) – see Figure 4. The coalition planning process can also be multi-staged. Such an approach requires substantial computational resources and fails in complex communities. For each single-staged CNP within a community of  $n$  agents, it is needed to send  $2(n+1)$  messages in the worst case.

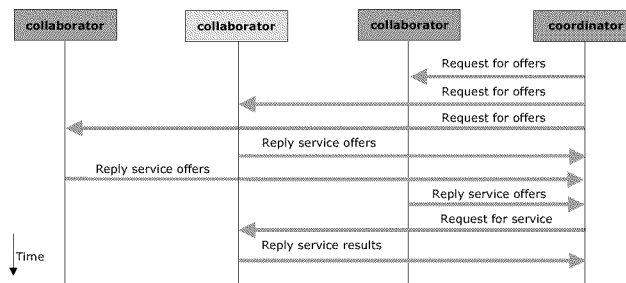


Figure 3 – Contraction based on a Single-staged Contract Net Protocol

At the same time many agents may not want to enter the CNP negotiation, as they wouldn't wish to undertake the risk of disclosing their private knowledge.

## 5.2 Acquaintance Model Contraction

The alternative communication strategy to CNP is based on exploitation of the agents' social knowledge. A coalition coordinator subscribes (by sending a `subscribe`-type of message) the potential collaborators for specific services they may want to exploit in the future. Upon a change in the collaborators' capabilities, they provide the coordinator with an update in the form of an `inform`-type of message. When the coordinator triggers the coalition formation phase, it parses the prepared service offers and selects the best collaborator(s) without any further negotiation. The coordinator sends a request, the collaborator updates its resources and confirms the contract. Any change in collaborator resources is advertised to all coordinators which subscribed the collaborator (see Figure 4).

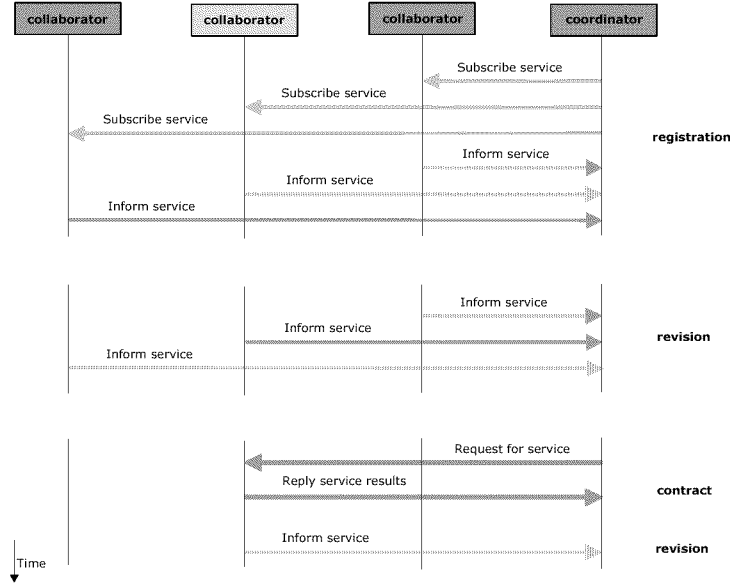


Figure 4 – Contraction based on Acquaintance Model exploitation

If there is a single event in the community  $\Theta$  that affects all the agents ( $n = |\Theta|$ ) and all the agents are mutually subscribed, then in the worst case there is  $(n(n-1))$  messages required for the social knowledge maintenance on this event. However, this is rarely the case. Agents never subscribe all each other (we could easily use a central communication component instead).

## 6 CPlanT Operation Lifecycle

The CPlanT multi-agent system operates in four separate phases: **registration** for agents' login/logout to/from the community, **alliance formation** for forming of alliances, **coalition formation** for finding a group of agents which can fulfill a well specified task and **team action planning** for resource allocation within the specific coalition. In the following, we will comment each of the phases.

### 6.1 Registration

Throughout the registration phase, a new-coming agent registers within the multi-agent community. This agent registers its public knowledge with the special central registration agent – the **facilitator**. Subsequently, the facilitator informs all the already existing agents about the new agent, and it also informs the new agent about all existing agents. Similarly, the agents can deregister with the facilitator. Any registered agent stores the public knowledge about all members of its total neighborhood  $\alpha(A)$  in the  $\text{Com-BB}(A)$  base of its acquaintance model. We have used the central communication unit – facilitator in the registration phase only. As the agents register just their public knowledge, we do not breach the requirements for confidentiality of the private information.

### 6.2 Alliance Formation

In this phase, which follows the registration process, the agents analyze the information they have about the members of the multi-agent system and attempt to form alliances. In principle, each agent is expected to compare its own private knowledge (i.e. alliance formation restrictions) with the public knowledge about the possible alliance members (i.e. type of an organization, its objectives, country of origin, etc.). Had the agent detected a possible future collaborator, the agent would propose joining the alliance. Throughout the negotiation process, the agent either

chooses the best alliance according its collaboration preferences of agents into already existing alliances. Failing to do so, an agent may start a new alliance by itself.

According to their preferences in Self-BB and community public knowledge in Com-BB, the agents carry out a selective contract net protocol process during this phase. The **quality of an alliance** is understood in terms of maximizing the individual agent's contribution to the alliance (i.e. covering the biggest amount of services that the other members of the alliance cannot implement). It is important to note that this process does not give us any guarantee for optimality of the alliance allocation. Each agent will join the most profitable alliance with respect to existing alliance configuration. With changing the order of agents' registration with the alliance, the formation algorithm will come up with different alliances.

### 6.3 Coalition Formation

In this phase, the agents group together not according to a similar mission objective, but they form coalitions with respect to a single, well-specified task that needs to be accomplished. Both, the CNP technique and the acquaintance model have been used in the coalition formation process. First, let us talk about the coalition formation process within a single alliance. The alliance members know the most of each other and are able to suggest a coalition that will very likely have foreseen properties. Whichever agent, member of an alliance, can face the role of the coordinator of the goal  $\tau$  implementation. The coordinator, who is to be set randomly in our implementation, parses its social neighborhood  $\mu(A)$  and detects the set of the most suitable collaborators (cooperation neighborhood)  $-\varepsilon(A, \tau)$ . Upon an approval from each of the suggested agents, the respective coalition  $\chi(\tau) = \varepsilon(A, \tau)$  is to be formed. Maintaining the agents' social neighborhood will save an important part of agent's interaction in the time of coalition formation. Agents will not need to broadcast a call for collaboration each time they will be required to accomplish a task. Instead, they will consult this pre-prepared knowledge and will contract the agent of which they knew it is the best to work with. The coordinator optimizes the **quality of a coalition** by seeking the coalitions that would contribute the most and in the shortest possible time.

As said in the previous, the agents' prefer not to form coalitions across alliances ( $\forall \tau: \varepsilon(A, \tau) \subseteq \mu(A)$ ). However sometimes an alliance fails to achieve a goal. The coordinator, who failed to form a coalition within one alliance, negotiates the proposal for collaboration by classical CNP with the agents from its total neighborhood  $\alpha(A_0)$ .

### 6.4 Team Action Planning

Once a coalition is formed, the agents share a joint commitment to achieve the goal  $\tau$ . Within this phase, a team of collaborative agents jointly creates a team action plan  $\pi(\tau)$ . The team action plan, that is a result of the coalition planning activity, is a joint commitment structure that defines exactly how each team member will contribute to achieving the shared goal (amount of resources, deadlines, etc.). The coordinator is supposed to (i) decompose a goal  $\tau$  into subtasks  $\{\tau_i\}$  and (ii) allocate the subtasks within the already formed coalition  $\chi(\tau)$ . There may be many achievable team action plans  $\pi(\tau)$ . The coordinator seeks for the cheapest or the fastest possible plan.

As there is no semi-private knowledge shared across the alliances, the agents from different alliances coordinate their activities by means of the contract net protocol. The intra-alliance team-action planning mechanism is not the pure acquaintance model contraction, where the team-action plan would result from the coalition leader deliberation process followed by a contract. All coalition members construct the precise team action plan collaboratively.

The collaborators advertise their services in the most informative while efficient form. We have suggested linear approximation of the discrete function that maps the delivery amount into due dates. Therefore the coordinator's acquaintance model stores the social knowledge that is imprecise, but very compact and efficient to parse. According to this social knowledge, the coordinator suggests the most optimal request decomposition and resource allocation  $-\pi(\tau)$  and transforms it into a contract proposal. This proposal is sent to the other coalition members, which reply with a specific collaboration proposal. However, the coordinator may find this proposal to be different than expected owing to the fact that the approximate information provided by the collaborator was far to imprecise. Instead of agreeing upon a joint commitment for this sub-optimal team action plan, the coordinator adapts the conflicting social knowledge and fires another round of negotiation. With each further negotiation stage the team action plan should be closer to the optimal team action plan. This process is to be iterated until there is no conflict in the expected capacity of the collaborators and the proposed delivery.

## 7 Implementation and Testing

### 7.1 Implementation

Testing correctness of the CPlanT required a well-defined, formal, but realistic enough scenario that can represent, model and initiate all aspects of agents' nontrivial behavior. The above specified principles and ideas have been

tested and implemented on a subset of the OOTW types of operations – humanitarian relief operations. For this purpose we designed and implemented a hypothetical humanitarian scenario Sufferterra representing a suffering island and several imaginary countries ready to help. The Sufferterra scenario was inspired by (Walker, 1999; Rathmell, 1999, Reece & Tate, 1998). The scenario knowledge has been encoded in XML and the computational model of the scenario has been implemented in Allegro Common Lisp.




Figure 5 – Sufferterra – subject of humanitarian operations

The R-Agents specify the physical arrangements of the geographical objects and the resources they provide. The problem specification does not distinguish the level of modeling granularity, i.e. each physical object may be implemented as an R-agent or several physical objects can make together an R-agent. For the testing purposes we have implemented a single R-Agent that represents the entire map of the area. The H-agents subscribe the R-Agent for specific information, by which these subscribers are informed about any change in physical arrangements of the relevant part of the map. There is a simple IN-Agent implemented as a part of the CPlanT community. Through one of the running instances of the IN-Agent, one can compose a “call-for-help” request and execute the coalition planning process. Such a request includes the type of disaster (“volcanic”, “hurricane”, “flood”, “earthquake”), the degree of disaster (1..9), location and the targeted H-Agent.

```
<city>
  <name> "Suffer Town" </name>
  <national-composition> "((christian 67) (muslim 18) (native 13) (other 2))"
</national-composition>
  <population> "50000" </population>
  <seaport>
    <ID> "1" </ID>
    <capacity> "25" </capacity>
    <material-hour> "200000" </material-hour>
  </seaport>
  <airport>
    <ID> "1" </ID>
    <capacity> "30" </capacity>
    <material-hour> "100000" </material-hour>
    <runway> "3000" </runway>
  </airport>
</city>
```

Figure 6 – Example of XML definition encoding of the ‘Suffer Town’ object

CPlanT has been successfully tested on the Sufferterra humanitarian relief scenario. The implementation is complemented by a visualizing meta-agent, which is implemented in Java. This meta-agent views logical structure of the system e.g. alliances, coalitions, team action plans and other properties of the community. There is a separate visualization for communication traffic monitoring. This component, that is not an agent, but rather a part of the multi-agent platform, serves mainly to debugging purposes. The community can be viewed and the requests can be sent from the web server via classical Internet browsers and from the WAP phones interface as well.

## 7.2 Experiments, Testing

Several different objectives were followed within the frame of the experiments: to evaluate the communication and computation requirements, quality of the solution provided and disclosure of private and semiprivate knowledge.

### Communication traffic

As stated in Section 5, an important part of the agent deliberation process has been decomposed into the inter-agent negotiation process. This is why we have concentrated our attention primarily to savings of the communication traffic in the entire system. The communication traffic has been observed in different architecture arrangements of

the community (e.g. different number of alliances) and for different complexity of the tasks sent to the community (e.g. different number of contracts). Having 20 agents we have experimented with the sample of all agents in one alliance, with agents clustered into 2, 4, 7 and 20 alliances. All the experiments have been carried out on the set of 19 measurements for each community arrangement. From the definition of the community lifecycle (see Section 6) follows that the latter case ( $\forall A: \mu(A)=\emptyset$ ) does not exploit any advantages of the acquaintance model contraction and the community behaves such as no social knowledge is administered and used. An important part of the communication traffic is carried out in the critical time – i.e. in the moment when the system is requested to provide a plan. By exploiting social knowledge that has been prepared in advance, we aimed at minimising communication traffic in this moment. The cost we have paid for this was the increased communication traffic in the idle times of the community. In the idle times, the agents are busy with maintaining the social knowledge stored in their acquaintance models. The communication traffic grows with the increasing number of alliances as each alliance member stores a more voluminous acquaintance model and it searches for a coalition by parsing the acquaintance model only.




Figure 7 – Communication traffic in communities with different number of alliances. The light bar depicts the maintenance messages, while the dark bar illustrates the overall communication in the system.

From the graph in Figure 7 we can see that with an increasing number of alliances (and a decreasing average number of alliance members) we reduce the communication requirements for maintenance of the model. The most of the communication in the critical time (the difference between dark and light bars in the graph) we save in the case of just one huge alliance. The optimal arrangement of the community was identified in the case of four alliances. However, it is not possible to define an optimal system structure because the agents cannot predict future tasks and the number of agents required for implementing these tasks. It is clear that for tasks requiring low number of agents, we will prefer small alliances while for the task requiring many agents, larger alliances will be preferred. The, the optimal size of a coalition is given by the nature of the tasks/goals under consideration.

### Evaluation of quality of the coalition

The evaluation of quality of the formed coalition is an important aspect in any coalition formation research. In Sufferterra scenario, there are two key attributes that influences the coalition value: (i) **success rate** – how many of the requested resources the coalition provides and (ii) **delivery time** – by when the coalition delivered the resources to the requestor. Experiments did not give any evidences to conclude any dependency between the success rate of the coalition and the used communication mechanism. However, with an increasing number of alliances, the overall delivery time is kept increasing due to additional costs of coordination among the alliances.

### Knowledge disclosure

The key challenge has been minimization of the private and semi-private knowledge disclosure. We have tried to measure both types of the information disclosures. Once the **private** information has been identified by another agent, this agent finds about the intent of the respective agent. As already noted, this very often happens when an alliance fails to plan all the requests and starts a contract net protocol process within members of the other alliances. Those, who will not be awarded the contract, know that the coordinator intends to operate in a mission and that it needs the resources requested.

The **semiprivate** information is disclosed in the same situation, when the possible collaborator proposes a service (as a reaction to a coordinator call for collaboration) that will not be accepted by the coordinator. In such a case, the coordinator finds out about the services the suggested collaborator can provide. Both the above mentioned cases are classified as strong knowledge disclosures (see Section 3.4). The weak knowledge disclosure happens in the registration phase within a single alliance and represents the amount of information that has become shared within the alliance.




Figure 8 : The graph on the left-hand side shows the dependence of the amount of private information disclosure in different architectures of the community. The graph on the right-hand side illustrates disclosure of the semi-private knowledge. The light bar depicts the weak and the dark bar strong knowledge disclosure.

As expected, the biggest disclosure of intents appears in the case of 20 alliances, as there is the highest CNP-based communication traffic among the alliances (see Figure 8). There is no weak disclosure once the agents are utterly independent (20 alliances). On the other hand, there is no strong semiprivate information disclosure in one alliance while the independent agents are starting to loose their semi-private information in the strong sense. It makes no implication to put together the strong and weak knowledge disclosures because of their different nature.

An interesting fact is that neither of the two extreme cases is the best for concealing the agents' private and semi-private knowledge. With one alliance, the semi-private knowledge becomes public while with no alliance each contract net protocol will reveal information about the contractors' intentions. It is rather difficult to find a good compromise in a number of alliances. What matters, is the probability that a request will not be fulfilled within one alliance and the coalition leader will have to subcontract other agents. Amount and structures of alliances in our domain emerge naturally according to the agents' private knowledge and other collaboration restrictions. Therefore it makes no sense to suggest an optimal number of alliances for a given community.

## 8 Relation to Coalition Planning Research

There has been a lot of work carried out in the area of coalition formation and coalition planning. It has been shown that finding the optimal coalition is an NP complete problem (Sandholm & Lesser, 1997). Researchers mainly suggest different negotiation strategies and analyze complexities of the coalition formation process (Shehory & Kraus, 1995). When a subject of optimization is the quality of the formed coalition, the agents usually act **collaboratively**. There have been published many of centralized planning mechanisms for coalition formation (Sandholm et. al., 1999). On the other hand, the **self-interested** agents maximize their own profit when participating in a coalition, no matter how well they will perform as a group. Many researchers analyzed properties of communities of self-interested agents such as their stability, worst-case profit, or payoff division among the agents (Li & Sycara, 2001). The domain we have investigated is partially of cooperative and self-interested type at the same time. Humanitarian aid providing agents tend to cooperate in the time of a crisis while they are self-interested and compete each other in a long-term horizon. Therefore, there was suggested a concept of alliances – collectives of agents that agreed to collaborate (to potentially form a coalition).

More importantly, the profit is very often the key optimization criterion when the agents optimize a coalition formation process (either collaboratively or competing each other). Besides the quality of the coalition, in the OOTW domain there are two (maybe more important) aspects to be taken into account. As forming an optimal coalition is a very complex problem, the **response time** becomes an important issue. Agents are limited in resources and a reasonably good answer, that is quickly provided, is very often much better than an optimal coalition found later (Steinmetz et. al, 1998; Sandholm & Lesser 1997). Practitioners would add that implementing a multi-agent system with a large number of agents, that are supposed to interact heavily, results in a **communication traffic overload** (Kaminka et. al., 2001). In our research we have tried to decompose the reasoning process and distribute it among the agents. While keeping the agents' deliberation process simple, we have concentrated our efforts on minimizing the communication interaction among the agents in order to suggest community structuring in a reasonable time. As the OOTW agents are also self-interested in certain way, they want to stay hidden in front of someone and advertise its collaborative capabilities to others. This is why we have to respect also the amount of **private information** to be disclosed. Therefore, we have also studied leaks of private information while forming the coalitions.

Research of the teamwork in a similar domain (evacuation scenarios) was reported in (Tambe, 1997). It was suggested to integrate the already existing software applications in the TEAMCORE wrapper agents. Unlike our acquaintance model that contains just social knowledge, the TEAMCORE wrapper agents also maintain domain specific team plans and the hierarchy of goals. Teams of agents share a team-oriented program, which is a joint knowledge structure that coordinates their activities. In CPlanT, there is no explicit team action plan distributed in

agents' acquaintance models. The structure of the coalitions and the team-action plan is a result of the inter-agent negotiation process. However, combination of both approaches where the agents' behavior is coordinated by a team-action plan that results from the agents' negotiation seems to be an interesting topic for further research.

Investigators approaching the problem from the game-theoretic point of view solve the problem of a higher complexity. Whereas in our case, there is a hierarchy structure for each task that is sent to the community and each task is coordinated by a single agent (the coordinator), in (Klusch et. al. 1997) all agents are equal. The agents autonomously analyze their own value. Through negotiations, they try to find out which coalition is the most profitable for them to join. This problem is inherently more complex and causes communication problems in complex communities. There will be several stages of negotiations needed as in many cases optimality of cooperation between two agents may not be reciprocal. In our case, we did not need to solve such a complex problem. On the other hand, in CPlanT we must optimize not only which coalition to join but also which services to provide to the coalition (e.g. team action planning). One may suggest that the game-theoretic approach could be used in the alliance formation phase of our algorithm (see Section 6.2). However, the agents join the system continuously, which makes it rather difficult to maintain the overall optimality of the distribution of alliances.

Besides the contract-net-protocol, there are other negotiation strategies based on classical auctioning mechanisms. While in combinatorial actions, the motivation of an agent is usually to make the biggest profit (or to contribute to a coalition in the best way), in our case, all the auctioneers and the bidding agents collaborate. A bidding agent tries to provide the auctioneer with such a bid that approximates in the best way the resources it can provide, and will help it to suggest the best possible resource allocation. In CPlanT, the agents also do not speculate about whom to work with. As we optimize the private information loss, collaboration within one alliance is always preferred. There is a potential of using the optimization for multiple auctioning mechanism for the team action planning within several overlapping coalitions (Anthony et. al. 2001).

## 9 Conclusions

The research described in this paper contributes to the coalition formation community by suggesting an alternative, knowledge based approach to the problem. Our research has been driven by the very specific domain of the OOTW. Apart from the classical contract net protocol techniques, we have used the communication strategy based on combination of three techniques: the centralized registration, the acquaintance models and the contract net protocol negotiations.

The agents in the community are organized into smaller, disjunctive groups called alliances. Each agent in the alliance is able to start the negotiation process to form a coalition and to develop a team action plan for a specific task either within the alliance or in collaboration with other alliances. Inside-alliance negotiations explore mainly the social knowledge stored in the acquaintance models, but the CNP technique is used as well (especially in the phase of the team action planning). The inter-alliance negotiations are based just on the CNP principles.

The general complexity of negotiations when forming a coalition in a MAS is of an exponentially explosive nature (Ketchpel 1993, Sandholm, 1995, Shehory, 1998). It has been shown that finding an optimal coalition is an NP complete problem when no specific constraints are imposed. In our case, the negotiation complexity of the coalition formation problem has been significantly reduced because:

- agents are organized into several disjunctive sets (alliances) and the most of coalitions are created just inside an alliance (reduced space of negotiations)
- the coalition leader within an alliance is set randomly (each coalition member has got the same coordination capacity and can manage the negotiation process), they don't compete for the role.
- within an alliance, the negotiation process explores the acquaintance models (social knowledge) in combination with the CNP technique and the pure CNP negotiations are used just in the case of the inter-alliance negotiations.

While the contract net protocol runs rather inefficiently, it keeps the agents from different alliances independent (they do not have to disclose their semi-private knowledge across alliances). This is why, the acquaintance-model based planning has been used exclusively within the alliances.

In our approach, we have not prioritized the requirement for the global coalition optimality, as this is not the main challenge in the OOTW planning. The main issue has been to develop an acceptable plan without forcing the agencies (agents) to make their private knowledge (namely intents and resources) public. This quite specific OOTW requirement enabled to reduce the complexity of the negotiation problem significantly. It has been measured that optimality of the coalition value slightly increases with the number of alliances (the role of the acquaintance model is getting smaller), while the problem complexity with a smaller number of socially knowledgeable alliances is significantly reduced.

## Acknowledgement

This research, was supported by AFOSR/European Office of Aerospace Research and Development under contract number F61775-00-WE043 as well as by the Grant No. LN00B096 of the Ministry of Education, Youth and Sports of the Czech Republic.

## References

- Anthony P., Hall W., Dang V. and Jennings N. R. (2001) "Autonomous agents for participating in multiple on-line auctions" *Proceedings of IJCAI Workshop on E-Business and the Intelligent Web*, Seattle, WA, 54-64
- Cao, W., Bian, C.-G., and Hartvigsen, G. (1997) Achieving Efficient Cooperation in a Multi-Agent System: The Twin-Base Modelling. In: *Co-operative Information Agents* (Kandzia, P., Klusch, M. eds.), LNAI No. 1202, Springer Verlag, Heidelberg, 210–221.
- Kaminka, G. A., Pynadath, D. V., and Tambe, M. (2001) "Monitoring Deployed Agent Teams", In *Proceedings of the Fifth International Conference on Autonomous Agents (Agents-2001)*, Montreal, Canada.
- Ketchpel, S. (1993) "Coalition Formation Among Autonomous Agents", *Proceedings of 5th European Workshop on Modeling Autonomous Agents in a Multi-Agent World*, Neuchatel, Switzerland, August 25-27
- Klusch, M., Contreras, J., Wu, F., Shehory, O. (1997) "Coalition Formation in a Power Transmission Planning Environment" *Proc. 2. Int. Conf. on Practical Applications of Multi-Agent Systems PAAM-97*, 21.-23.4, London
- Li C. and Sycara K. (2001) "Algorithms for combinatorial coalition formation and payoff division in an electronic marketplace" *Tech. report CMU-RI-TR-01-33*, Carnegie Mellon University
- Mařík V., Pěchouček M., Štěpánková O., Lažanský J. (2000) "ProPlanT: Multi-Agent System for Production Planning", *International Journal of Applied Artificial Intelligence*. 2000, vol. 14, no. 7, p. 727-760
- Mařík V., Pěchouček M., Štěpánková O. (2001) "Social Knowledge in Multi-Agent Systems", In (M. Luck et. al, eds.) *Multi-Agent Systems and Applications*, LNAI 2086, Springer-Verlag, Heidelberg, 2001, pp.211-245
- Pěchouček, M., Norrie, D. (2000) "Knowledge Structures for Reflective Multi-Agent Systems: On reasoning about other agents", registered as Report Number 538, Department of Mechanical and Manufacturing Engineering, University of Calgary, Alberta, Canada,
- Pěchouček M., Mařík V., Štěpánková O. (2001) "Towards Reducing Communication Traffic In Multi-Agent Systems", *Special Issue on Virtual Organizations and E-Commerce Applications*, *Journal of Applied System Studies*, Cambridge International Science Publishing, ISSN 1466-7738, Cambridge, UK, Spring, 2001.
- Rathmell R. A. (1999) "A Coalition Force Scenario 'Binni – Gateway to the Golden Bowl of Africa'", *Defense Evaluation Research Agency*
- Reece, G. A. and Tate, A. (1993) "The Pacifica NEO Scenario". Technical Report ARPA-RL/O-Plan2/TR/3. Artificial Intelligence Applications Institute, University of Edinburgh, Scotland
- Sandholm, T. and Lesser, V. (1997) Coalitions among Computationally Bounded Agents. *Artificial Intelligence* 94(1), 99-137, Special issue on Economic Principles of Multi-Agent Systems.
- Sandholm, T., Larson, K., Andersson, M., Shehory, O., and Tohmé, F. (1999). Coalition Structure Generation with Worst Case Guarantees. *Artificial Intelligence*, 111(1-2), 209-238.
- Shehory O. and Kraus S. (1995) Coalition formation among autonomous agents: Strategies and complexity, *Lecture Notes in Artificial Intelligence* no. 957, *From Reaction to Cognition*, (Castelfranchi & Muller, eds.), pp 57-72,
- R. Smith (1980) The contract net protocol: High-level communication and control in distributed problem solver. *IEEE Transactions on Computers*, 29(12):1104--1113
- Steinmetz, J. Collins, M. Gini, and B. Mobasher (1998) "An Efficient Anytime Algorithm For Multiple-Component Bid Selection in Automated Contracting" , In *Agent Mediated Electronic Trading*, LNAI1571, Springer, pp 105-125
- Tambe M. (1997) Towards flexible teamwork. *Journal of Artificial Intelligence Research*, 7:83 – 124
- Walker, E. C. T. (1999) Panel Report: Coalition Planning for Operations Other Than War, Workshop at AIAI, Edinburgh
- Wooldridge M. (2000) *Reasoning about rational agents*, The MIT Press, Cambridge, MA,